

AFRL-IF-RS-TR-2005-5
Final Technical Report
January 2005



PROVIDING SURVIVABLE REAL-TIME COMMUNICATION SERVICE FOR DISTRIBUTED MISSION CRITICAL SYSTEMS

Texas A&M University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. H564 & J034

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2005-5 has been reviewed and is approved for publication

APPROVED: /s/

ALAN J. AKINS
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 074-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE JANUARY 2005	3. REPORT TYPE AND DATES COVERED Final Jun 99 – May 04	
4. TITLE AND SUBTITLE PROVIDING SURVIVABLE REAL-TIME COMMUNICATION SERVICE FOR DISTRIBUTED MISSION CRITICAL SYSTEMS			5. FUNDING NUMBERS C - F30602-99-1-0531 PE - 62301E PR - H564 TA - 10 WU - 01	
6. AUTHOR(S) Wei Zhao, Riccardo Bettati, and Nitin Vaidya				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Texas A&M University H. Bright Build, CS Department College Station Texas 77845			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFGA 3701 North Fairfax Drive 525 Brooks Road Arlington Virginia 22203-1714 Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2005-5	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Alan J. Akins/IFGA/(315) 330-1869/ Alan.Akins@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This document is the final report for Providing Survivable Real-Time Communication Service for Distributed Mission Critical Systems, a Texas A&M project funded through the DARPA Fault Tolerant Networks Program. In this project, we have developed techniques for survivable real-time services for mission critical systems. In particular, we incorporated real-time traffic modeling techniques into the security service to enhance both system security and real-time capabilities in an adverse environment. This is the key element that has made our work innovative and unique. We utilized a combination of traffic stuffing based on traffic-modeling and intrusion detection/monitoring as highly effective countermeasures against both traffic analysis and denial of service forms of attack in both wired and wireless networks. The threat of traffic analysis is particularly critical in wireless environments, where the operational mode of the critical applications can be easily inferred. The project has been successful. We developed a system called NetCamo which integrated the technologies we developed into a deployable form. NetCamo was successfully integrated into the Navy's Hi-Per-D system. Further, NetCamo was transferred to an industrial company for deployment in both DoD and commercial domains.				
14. SUBJECT TERMS Fault Tolerant Networks, Network Security, Data Camouflage, Traffic Analysis Attack, Traffic Modeling, Denial Of Service, Intrusion Detection				15. NUMBER OF PAGES 18
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

TABLE OF CONTENTS

1. Introduction.....	1
1.1. Objective.....	1
1.2. Challenges	1
1.3. Proposed Solution.....	2
1.4. Contribution.....	3
1.5. Report Organization	3
2. Technical Background	3
3. Technical Approaches.....	6
3.1. Countermeasure Algorithms.....	6
3.2. NetCamo System	7
4. Results and Discussions.....	9
5. Conclusions and Future Work	12
6. References.....	12

LIST OF FIGURES

Figure 2-1: Observed and Derived Traffic.....	5
Figure 3-1: NetCamo Base Architecture.....	8

1. Introduction

This document is the final report for *Providing Survivable Real-Time Communication Service for Distributed Mission Critical Systems*, a Texas A&M project funded through the DARPA Fault Tolerant Networks program.

1.1. Objective

In this project, we developed techniques for survivable real-time services for mission critical systems. We incorporated *real-time traffic modeling techniques* into the security services to enhance both system security *and* real-time capabilities in an adverse environment. This is the key element that makes our work innovative and unique. In particular, our objective was to develop techniques for survivable real-time services for mission critical systems. Specifically, we focused on:

- 1) Developing technologies that effectively shield network resources and payload connections from passive attacks and that support rapid recognition of active attacks;
- 2) Developing tool sets that help to build mission critical systems that exhibit inherent survivability properties, i.e., the ability to continue real-time operation in the face of attacks that are partially successful.

1.2. Challenges

Traditionally, encryption has played an important role in network security. However, it is a misconception that to secure a network, one only needs to encrypt the traffic. With increasing amounts of traffic being encrypted and its contents therefore being beyond the reach of effective cryptanalysis, attention is shifting towards *traffic analysis*, and the prevention thereof. Traffic analysis is a security attack where an intruder observes network traffic in order to infer sensitive information about the applications and/or the underlying system. This form of attack is harmful because significant information about operational modes can be inferred by appropriately monitoring the pattern of traffic. It can, for example, uncover the location of command centers, determine the state of alertness of various units, or detect covert information

flows to or from apparently non-involved parties. In addition, effective traffic analysis is well known to greatly help the cryptanalysis efforts. It is therefore important to develop a means to render traffic analysis efforts ineffective. Traffic analysis can be prevented by *camouflaging* the payload traffic, i.e., manipulating the traffic so that, to an observer, its pattern is not related to the operational status of applications. In order to achieve this, an integration of the following measures should be used:

- Traffic Padding. Additional packets (called padding packet) may need to be properly inserted into payload packet streams to camouflage them.
- Traffic Re-Routing. Usually, packets from one host to another are sent via one fixed path. In order to prevent traffic analysis, a stream of traffic between two hosts may need to be re-routed through multiple paths in order to camouflage the traffic.

The challenge of this study was to deal with the problem of preventing traffic analysis in the context of mission critical system where the worst-case delay of payload packets needs to be guaranteed. This is not possible when the network is indiscriminately flooded by padding traffic.

1.3. Proposed Solution

We used *traffic-modeling based traffic stuffing* as a highly effective countermeasure against the passive attack in both wired and wireless networks. The threat of traffic analysis is particularly critical in wireless environments, where the operational mode of the critical applications can be easily inferred. In this project, we applied our traffic modeling techniques in network security. In particular, for *traffic analysis* attacks, we use traffic stuffing algorithms that can *effectively mask the actual operational modes of mission critical applications* without compromising the guaranteed quality of service (QoS). This is achieved by using the traffic modeling theory to precisely flood the network at the right time and the right place.

1.4. Contribution

The project produced the following specific results:

- We developed and analyzed security countermeasures based on traffic modeling.
- We analyzed the impact of wireless environments in terms of passive attacks and their countermeasures.
- We implemented a tool set (called NetCamo) for integrated security and real-time services.

In terms of technology transfer and deployment, we performed demonstrations at DARPA PI meetings, participated in the DARPA Fault Tolerant Networks (FTN) experimental tests led by BBN, integrated NetCamo with the Navy's HiPer-D system, and transferred NetCamo to an industrial company for future deployment in both DoD and commercial domains.

1.5. Report Organization

There are five sections in this report. Section 2 covers the technology background, including discussion of real-time modeling theory. Section 3 discusses our technical approach and considers attacker capabilities along with defender strategies. It also outlines the fundamental concepts of the NetCamo system. Section 4 reports the main findings of this project while Section 5 covers conclusions and recommendations for future work.

2. Technical Background

At Texas A&M University, prior to this effort we developed various techniques to provide delay-guaranteed communication services in high performance networks including token ring, FDDI, ATM, wormhole networks, etc. In these previous research and development projects, we discovered that traffic modeling plays a critical role in delay-guaranteed communications. Only when we properly model the traffic in the network, can we efficiently derive tight worst-case delay bounds and hence provide delay guarantees to mission critical applications.

Traditionally, stochastic models have been used to characterize network traffic [MZ94]. However, such models only provide insight into the average performance of the network, which is not adequate for mission critical real-time systems. To provide delay-guaranteed communication, we seek mathematical tools that can describe bounds on the behavior of the systems. In particular, we developed [RKZ95, FLR96, SCZ97] the following mathematical functions to characterize the traffic from a source at an arbitrary point of the network:

$$F(I) = \max_t(A(t+I) - A(t)) \quad (1-1)$$

and

$$f(I) = \min_t(A(t+I) - A(t)) \quad (1-2)$$

where $A(t)$ is the number of bits that have arrived by time t . $A(t+I) - A(t)$, thus describes the number of bits that arrived during the time interval $(t, t+I)$. Consequently, $F(I)$ and $f(I)$ define the maximum and minimum numbers of bits that arrive during *any* interval of length I , respectively. In the case that the traffic is constant-rated, then


$$F(I) = f(I) = a I \quad (1-3)$$

where a is the rate.

Using these functions, we developed various methods to predict the worst-case delays of messages, as required by our previous real-time communication projects [ACZ92, SCZ97]. The following is a brief overview the major relevant results:

1. These traffic-modeling functions can describe a wide variety of input traffic sources [SCZ97]. Different sources generate different kinds of traffic. For example, the traffic generated by a multi-media application may be very different from that generated by the process for a monitoring and control application.
2. Cost effective representation of these traffic functions can be obtained [SDG98]. For example, we proposed and validated a six-point representation. That is, by using six pairs of function values, a maximum (minimum) traffic function can be effectively approximated without losing much information. This is critical in real-time manipulation during the intrusion detection process. Figure 2-1 is an example of the observed and derived maximum traffic function. It is shown in the

figure that using our six-point approximation model, the derived function is extremely close to the observed function.



Data source: A SUN workstation transmitting a message of 16Mbits per second
Network: ATM with line speed of 155Mbits/S

Figure 2-1: Observed and Derived $\Gamma(I) = F(I)/I$

Figure 2-1: Observed and Derived Traffic

3. These traffic functions can be analytically derived [FLR96, NSZ97]. Let F^{in} , F^{out} , f^{in} , and f^{out} be the maximum and minimum traffic functions at the input and output of a router, respectively. Then, the following formulas can be proven:

$$F^{out}(I) = F^{in}(I + d) \quad (1-4)$$

and

$$f^{out}(I) = f^{in}(I - d) \quad (1-5)$$

where d is the worst case delay at the router. The value for d depends on the scheduling methodology used in the switch and can be obtained by various analytical techniques. Formulas (1-4) and (1-5) imply that the traffic at the output

of a router, modeled by F^{out} and f^{out} , can be derived from the traffic at its input (i.e., the output of the previous router). Consequently, traffic functions of a flow along its entire path can be derived once the source traffic is specified.

These observations, especially the third one, are particularly important for security services. Note that (1-4) and (1-5) suggest that, except for being (left or right) shifted, the traffic functions are almost identical to corresponding functions from the source host. The implications here are two fold:

- a) These traffic functions carry signature-like information from the source. Indeed, they characterize the behavior of the source in the various burst regions (cell, packet, and message). Since different hardware and software configurations directly affect the rate at which traffic leaves the sender, our traffic functions can be used to uniquely identify a traffic flow. This feature should be explored in intrusion detection and suppression.
- b) Furthermore, $F(I)$ provides the upper bound on traffic. A countermeasure of traffic analysis can use this information to effectively distribute “stuffing” traffic in order to mask the real operational mode. If flooding is combined with traffic modeling, no particular tagging is necessary to detect stuffing traffic, as it can be identified based on the real traffic function for the current operational mode.

In this project, we applied these techniques in building our survivable real-time communication networks.

3. Technical Approaches

3.1. Countermeasure Algorithms

Let us consider a simplified model first. Consider a network that consists of n nodes, $N_1, N_2 \dots N_n$. The application system may run in one of several operational modes (say, $M_1, M_2 \dots M_m$). Each mode is characterized by a set of communication connections that are active in the mode. A mode typically reflects an operational status of the underlying application. A connection is described by parameters: source node, destination node, input traffic functions ($F^{in}(\cdot)$ and $f^{in}(\cdot)$), and deadline. $F^{in}(\cdot)$ and $f^{in}(\cdot)$ are functions of the maximum and minimum traffic transmitted by the source,

respectively. If a packet is transmitted by the source at time t , it must be received by the destination node by $t + D$ where D is the connection's deadline.

Under this simple network model, the problem of countermeasure for passive attacks is to properly schedule the stuffing messages over the network such that an enemy cannot detect what the current operational mode is. Here we assume that the enemy can listen to (anywhere within) the network and even examine the packet header.

One simple way to mask the operational mode is to make all the operational modes appears the same. Technically speaking, this means that the aggregated traffic on each link should appear as constant rate traffic. That is, in terms of our traffic modeling terminology, for every link j ,

$$\sum_i F_{i,j}(I) + S_j(I) = \alpha I \quad (3-1)$$

where $F_{i,j}(\cdot)$ is the traffic function for the traffic over link j from connection i . $S_j(\cdot)$ is the traffic function for stuffing messages over link j and α is the desired constant rate that is invariant in all the links, for all the operational modes. Recall that $F_{i,j}(\cdot)$ can be derived by our real-time traffic modeling theory. Thus, we solve (3-1) to obtain $S_j(\cdot)$, which will consequently be used by the node that transmits to link j in order to correctly generate stuffing messages. By doing so, the operational mode is successfully masked. No matter what the current operational mode is, the enemy will observe the same traffic pattern. The above approach is then extended to the cases of multiple cover modes and the wireless domain.

3.2. NetCamo System

A critical part of this proposed project was to fully develop and implement a deployable network management and control system, called NetCamo. In this section, we will give an overview of the architecture of NetCamo and elaborate on its implementation.

An overview of the architecture of NetCamo is illustrated in Figure 3-1. The NetCamo system consists of two major components: the NetCamo *Network Controller* and the NetCamo *Host Controller*. The NetCamo Network Controller contains the NetCamo *Traffic Manager*, the *Router Agent* and the *Host Agent*. The NetCamo Traffic

Applications can call the NetCamo API to establish or teardown a connection. The connection admission or release request is submitted to a Host Agent and then to the NetCamo Traffic Manager for admission control. If this connection can be accepted, the NetCamo Traffic Manager notifies all the related NetCamo Host Controllers about the new connection and returns to the application with accept or reject information. Particular attention was paid to compatibility with existing establishment and resource allocation protocols for easy integration with existing application bases – we call this *upward compatibility*. For example, NetCamo provides a distributed H.323 Gatekeeper interface, which allows applications that are H.323 compliant (e.g., Microsoft NetMeeting) to be transparently deployed over NetCamo. Similarly, NetCamo can be easily integrated with other systems, such as RSVP or Honeywell’s RTARM resource allocation protocol.

Since its deployment, based on the feedback from our customers in both DoD and commercial organizations, several variations of the NetCamo system have been made to customize it for the particular need. For example, to meet the requirements of the Navy’s HiPer-D system, we modified the implementation plan and made the host traffic a standalone unit, rather than a middleware implementation that would have been inserted in the application hosts.

4. Results and Discussions

The NetCamo system prevents traffic analysis in systems with real-time requirements. This project has developed a complete implementation of the NetCamo system, starting from the general approach and design, to evaluation and deployment. Integrated support for both security and real-time requirements is becoming necessary for computer networks that support mission critical applications. This study focused on how to integrate both the prevention of traffic analysis and guarantees for worst-case delays in a network. We proposed and analyzed techniques that efficiently camouflage network traffic and correctly plan and schedule the transmission of payload traffic so that both security and real-time requirements are met.

We developed the methodologies used in the NetCamo system to prevent traffic analysis and guarantee message deadlines. We took an integrated approach and realized

it in three system phases: 1) At the system configuration phase, correct camouflage traffic parameters are derived in accordance with various constraints of the system. Delays along individual communication links are also analyzed. 2) At the admission control phase, a traffic plan is generated for the newly arrived connection. The traffic plan specifies how to distribute the traffic for the connection into several routers so that both camouflaging and delay requirements are met. 3) At the run-time phase, traffic control is applied at a low layer by sending additional dummy traffic on communication paths to compensate for fluctuations in traffic at the connection level so that the real traffic is consistent with the desired camouflaged traffic pattern.

Comprehensive performance evaluation has been carried out. The performance evaluation showed that the NetCamo system is effective and efficient. In NetCamo, the error between targeted time invariant traffic and real camouflaged traffic can be as small as one percent, and at the same time the probability of a guaranteed real-time connection is still relatively high. These facts indicate that with our innovative camouflage and traffic planning technologies, the NetCamo system can effectively camouflage the network traffic without compromising real-time capability. The reader is referred to <http://netcamo.cs.tamu.edu> for the latest developments on the NetCamo system.

Given the initial success of the NetCamo system, several extensions were investigated. We considered allowing the network to have multiple modes of camouflaged traffic patterns. This allows the network to mislead an intruder more effectively. For example, the camouflaged traffic could appear as if it were in a “peace time” status when actually an alert status is in progress. Conversely, when the network is really in peace time operation, we may want it to appear to be in an alert status. Another extension allowed the NetCamo network controller to be decentralized in order to improve its efficiency when working in a large distributed environment.

The NetCamo system was successfully integrated with DoD’s (the Navy’s) High Performance Distributed Computing system (HiPer-D), a joint effort between the Aegis shipbuilding program and several DARPA information technology programs. HiPer-D defined a new distributed real-time computing architecture for shipboard use with dynamic resource management that allows large-scale, real-time systems to dynamically

reconfigure themselves to adapt to varying environments, changing mission demands, and current resource availability. As pointed out earlier, to meet the requirements of the HiPer-D system, the implementation plan was modified to make the host traffic controller a standalone unit, rather than middleware inserted into the application hosts.

We also worked with Tri-Pacific Software to produce a deployable version of NetCamo at industrial quality. The NetCamo application, when fully developed and integrated with other COTS security software, will add an additional layer of security to network infrastructures that the market is poised to accept. Furthermore, large middleware-dependent enterprise system management (ESM) companies are just beginning to increase their focus on this growing market segment by developing new products and partnerships to address the security needs of current and potential customers. The additional denial of service protection and network management features included in NetCamo will prove valuable to any organization with a major network or a virtual private network (VPN). However, the majority of clients initially interested in NetCamo will likely consist of large companies with private networks, VPN's, network service providers, and governmental organizations. As noted by Gartner Vice-President John Pescatore, “ an ISP or Internet data center, or one of the large backbone centers like Sprint or AT&T ...need to implement denial of service protection.” Also, “large enterprises like state governments that act as ISPs for their agencies are also likely to be shopping for DoS protection.” (Security News & Analysis, Michael S. Mimoso, June 22, 2001).

The NetCamo system was also deployed at the Information Technology and Operations Center (ITOC) at the United States Military Academy. The system helped development of the curriculum and was used in the annual Cyber Defense Exercise. In addition, we also carried out collaboration with several academic institutions, including the University of Texas at Dallas, the University of Dallas, and the University of Texas at San Antonio.

Our work is the first to address both problems of preventing traffic analysis and guaranteeing worst case delays in an integrated manner. NetCamo is the first *implemented* system that can offer this kind of secure real-time communication service in an open IP network. The NetCamo system effectively camouflages the payload traffic pattern while traditional encryption technology camouflages the payload content.

5. Conclusions and Future Work

The main focus of this project was the development of the NetCamo system, including the general approach, design, implementation, evaluation, and deployment. NetCamo provides integrated support for both network security and real-time operations which are key features in computer networks that support mission critical applications. In general, our techniques efficiently camouflage network traffic and correctly plan and schedule the transmission of payload traffic so that both security and real-time requirements are met. NetCamo achieves very high levels of camouflaging without compromising real-time requirements [JVZ05, GFB04].

There are many other components that may need to be camouflaged in a communication system, including location and type of hosts, network topology, etc. We are studying various issues related to camouflaging these and other components that may be of interest to an intruder.

6. References

- [ACZ92] G. Agrawal, B. Chen, Wei Zhao, and S. Davari, Guaranteeing Synchronous Message Deadlines in High Speed Token Ring Networks with Time Token Protocol, Winner of the Best Paper Award. In *Proceedings of IEEE International Conference on Distributed Computing Systems*, pp. 468-475, June 1992.
- [FGB02] X. Fu, Y. Guan, R. Bettati, and Wei Zhao, Hiding Role Assignment for Mission Critical Collaborative Systems, in *Quality and Reliability Engineering International, Special Issue on Computer Network Security*, Volume 18, Issue 3, 2002, pp201- pp216.
- [FGB03] X. Fu, B. Graham, R. Bettati, and Wei Zhao, Active Traffic Analysis Attacks and Countermeasures, in *Proceedings of the 2nd IEEE International Conference on Computer Networks and Mobile Computing (ICCNMC)*, Oct 2003

- [FGB03b] Xinwen Fu, Bryan Graham, Riccardo Bettati and Wei Zhao, Analytical and Empirical Analysis of Countermeasures to Traffic Analysis Attacks, in *Proceedings of International Conference on Parallel Processing (ICPP)*, Oct. 2003.
- [FGX04] X. Fu, B. Graham, D. Xuan, R. Bettati and Wei Zhao, Empirical and Theoretical Evaluation of Active Probing Attacks and Their Countermeasures, in *Proceedings of the 6th Information Hiding Workshop*, May 2004
- [FLR96] F. Feng, C. Li, A. Raha, S. Yu and Wei Zhao, Modeling and Regulation of Host Traffic in ATM Networks for Hard Real-Time Applications, in *Proceedings of IEEE Conference on Local Computer Networks*, Oct. 1996.
- [GFS01] Y. Guan, X. Fu, D. Xuan, P. Shenoy, R. Bettati, and Wei Zhao, NetCamo: Camouflaging Network Traffic for QoS-Guaranteed Mission Critical Applications, in *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 31, No. 4, July 2001.
- [GFB04] Y. Guan, X. Fu, R. Bettati, and Wei Zhao, A Quantitative Analysis of Anonymous Communications, in *IEEE Transactions on Reliability*, Vol 53, No. 1, March 2004, pp 103 – 115.
- [JVZ05] S. Jiang, N. Vaidya, and Wei Zhao, A Mix Route Algorithm for Mix-Net in Wireless Ad Hoc Networks, in *Proceedings of the 1st IEEE International Conference on Mobile Ad-hoc and Sensor Systems*, Oct 2005.
- [MZ94] N. Malcolm and Wei Zhao, Hard Real-Time Communication in Multiple-Access Networks, *Journal of Real-Time Systems*, Vol. 8, No. 1, pp. 35 - 77, Jan. 1994.
- [NSZ97] J. Ng, S. Song, and Wei Zhao, Integrated Delay Analysis of Regulated ATM Switc", in *Proceedings of IEEE Real-Time Systems Symposium (RTSS)*, Dec. 1997.

- [RKZ95] A. Raha, S. Kamat, W. Zhao, "Guaranteeing End-to-End Deadlines in ATM Network", in *Proceedings of IEEE International Conference on Distributed Computing Systems*, May 1995.
- [SCZ97] A. Sahoo, B. Chen, and Wei Zhao, "Connection-Oriented Communications for Real-Time Applications in FDDI-ATM-FDDI Heterogeneous Networks", in *Proceedings of IEEE International. Conference on Distributed Computing Systems (ICDCS)*, May 1997.
- [SDG98] A. Sahoo, B. Devalla, Y. Guan, R. Bettati, and Wei Zhao, "Adaptive connection management for mission critical applications over ATM Networks", Winner of the Best Paper Award, in *Proceedings of IEEE Aerospace and Electronics Conference*, July 1998.
- [ZFG05] Y. Zhu, X. Fu, B. Graham, R. Bettati and Wei Zhao, On Flow Correlation Attacks and Countermeasures in Mix Networks, in *Proceedings of the Workshop on Privacy Enhancing Technologies (PET2004)*, Toronto, Canada, May 26-28, 2004.